

09/236753

#3

UK Patent Application GB 2 284 121 A

(43) Date of A Publication 24.05.1995

(21) Application No 9421658.7

(22) Date of Filing 27.10.1994

(30) Priority Data

(31) 107658

(32) 18.11.1993

(33) IL

(51) INT CL⁶
H04N 7/26

(52) UK CL (Edition N)
H4F FD12X FD3 FD30K FRX

(56) Documents Cited
WO 94/23385 A2 WO 93/17524 A1

(58) Field of Search
UK CL (Edition N) H4F FRT FRW FRX
INT CL⁶ H04N 7/26
Online databases : wpi

(71) Applicant(s)

State of Israel - Ministry of Defence

(Incorporated in Israel)

PO Box 2250, Armament Development Authority,
Rafael, Haifa 31 021, Israel

(72) Inventor(s)

Ido Bettess
Itzhak Florentin

(74) Agent and/or Address for Service

Urquhart-Dykes & Lord
3rd Floor, Victoria Suite, Northern Assurance
Buildings, Albert Square, MANCHESTER, M2 4DN,
United Kingdom

(54) Data compression using wavelet transformation; tree pruning

(57) Data, such as image data, to be compressed first undergoes wavelet transformation to produce a plurality of layers (H1 to H5) each having a different resolution. Each lowest resolution datapoint (layer H5) may be considered as having tree comprising a plurality of descendent datapoints at higher resolutions (H2 to H5); the lowest resolution datapoints may be considered as forming a forest of trees. To compress the data branches of trees having datapoints with low information content (ie below a threshold) are pruned from the tree structure and only the remaining data - figure 5B(i) - transmitted or stored. Datapoints below the threshold may be set to zero before pruning; zero-crossings may also be pruned from the tree.

FIG.5A





FIG.5B



GB 2 284 121 A




FIG. 1 A
PRIOR ART




FIG. 1B
PRIOR ART





FIG. 1C
PRIOR ART

4/13

FIG. 2
PRIOR ART







FIG.3

6/13

FIG. 4






FIG.5A




FIG.5B

FIG. 6




FIG.7





FIG.8/1






FIG.8/2




FIG.10A




FIG.10B




FIG.9




FIG.11

A SYSTEM FOR COMPACTION AND RECONSTRUCTION OF WAVELET DATA
FIELD OF THE INVENTION

5 The present invention relates to data processing systems generally and to systems utilizing wavelet transforms in particular.

BACKGROUND OF THE INVENTION

10 Data processing typically involves transforming the data from one space to another. The best known transformation is the Fourier transformation which transforms a signal from a time space to a frequency space, indicating thereby the frequency components of the input signal.

15 Recently, wavelet transforms have been developed where each wavelet transform has a different curving "wavelet" shape and the result of the transformation indicates to what extent the input signal, a digital signal, is formed of a combination of the different wavelets. The article "Wavelets and Signal Processing", by Olivier Rioul and Martin Vetterli, published in IEEE SP Magazine, October 1991, pp. 14 - 38, describes prior art wavelet transform operations and is incorporated herein by reference.

20 Typically, wavelet transforms are implemented in a multi-resolution manner. The prior art multi-resolution method and exemplary output, for a one-dimensional signal, are illustrated in Figs. 1A and 2, respectively, to which reference is now made.

25 Shown in Fig. 1A are three "wavelet transform units" 10a - 10c. Each wavelet transform unit 10a - 10c is formed of its respective wavelet low pass filter (WLPF) 12a - 12c, wavelet high pass filter (WHPF) 14a - 14c and samplers 16a - 16c. The wavelet low and high pass filters 12a - 12c and 14a - 14c break the signal into its high and low wavelet frequency components and the samplers 16a - 16c sample the output of each filter, thereby changing the resolution of the signal. For one-dimensional input signals, the samplers 16a - 16c typically take every second datapoint, thereby reducing the number of points by two with each level. Other sampling rates can also be implemented.

30 Each transform unit 10a - 10c operates on a different resolution signal. Therefore, the elements of each transform unit are labeled to indicate the resolution level or

iteration, 1 - 3, at which they operate, wherein the highest level has the lowest resolution (least number of points). The first transform unit 10a operates on the input signal and the other transform units 10b - 10c operate on the low frequency output signals of the transform unit 10a - 10b, respectively, of the level previous to it. The signals of Fig. 1 are labeled to indicate the level and whether they are the low (L) or high (H) frequency output of the level.

The output of a wavelet transform operation comprises the following signals, the high frequency output signals H1 - H3 and the low frequency output signal L3 of the last level. The low frequency output signal of the highest level, in this example signal L3, is herein called the "core" signal.

An exemplary input signal, labeled 18, and its corresponding three high frequency output signals are shown in Fig. 2. As can be seen, the resolution is reduced at each level. Furthermore, the activity of the lower resolution signals is near at the inflection point of the input signal 18.

If the input signal 18 is an image or some other two-dimensional signal, the wavelet low and high pass filters can either be separable or non-separable. The operations for separable and non-separable filters are shown in Figs. 1B and 1C, respectively, to which reference is now made.

The wavelet transform units 10 for separable filtering include low and high pass wavelet filters 11a and 11b (Fig. 1B) and samplers 13a for rows and low and high pass wavelet filters 11c and 11d and samplers 13b for columns. The rows of the input signal first pass through the low and high pass wavelet filters 11a and 11b and are sampled by the samplers 13a, thereby reducing the number of pixels in each row by two. The columns of the row filtered signals then pass through the low and high pass wavelet filters 11c and 11d and are sampled by the samplers 13b, thereby reducing the number of pixels in the columns by two. Alternatively, the columns can be operated on first, followed by the rows.

The separable filtering method produces three high frequency signals Ha, Hb, Hc and one low frequency signal L at each level. Thus, the first level produces the signals H1a, H1b, H1c and L1. Each signal has one-fourth the number of pixels as the input signal. The low frequency signal L1 is then provided as the input signal to the next wavelet transform unit 10.

The wavelet transform units 10 for non-separable filtering operate differently for even- and odd-numbered levels, as illustrated in Fig. 1C. For odd levels, the input signal

5 is operated on by non-separable low and high pass wavelet filters 15a and 15b, respectively, and quincunx odd samplers 17a. Odd samplers 17a perform "quincunx" subsampling which removes every even datapoint in odd numbered rows and every odd datapoint in even numbered rows, leaving a checkerboard-like output image, illustrated in part in Fig. 10B. In
order to maintain the image, the datapoints are not actually removed; instead, they just are
not operated upon by the following even level.

10 For even levels, the wavelet transform unit 10 operates on the checkerboard-like image with non-separable low and high pass wavelet filters 15c and 15d, respectively, and quincunx even samplers 17b. Even samplers 17b also perform quincunx sampling on the checkerboard-like image, resulting in a more normal looking image, shown in part in Fig. 10A, which is the input to the odd-level wavelet transform unit. The image produced by even samplers 17b has one-half the number of elements as the input image to the odd samplers 17a.

15 The article "Predictive Interscale Image Coding Using Vector Quantization", by M. Antonini et al and published in Signal Processing V:Theories and Applications, L. Torres, E. Masgrau and M.A. Lagunas, editors, Elsevier Science Publishers, B.V. 1990, pp. 1091 - 1094, describes an image compression method using biorthogonal wavelet transforms and an interscale prediction scheme. The prediction scheme predicts the location and amplitude of edges in an image at one resolution based on their locations and amplitudes at
20 the next highest resolution.

The article "Image Coding using Lattice Vector Quantization of Wavelet Coefficients", by M. Antonini et al., ICASSP 91, pp. 2273 - 2276 describes a compression method for wavelet transform data.

25 U.S. Patent 5,014,134 to Lawton et al. describes an image compression and decompression system which utilizes wavelet transforms.

SUMMARY OF THE PRESENT INVENTION

It is an object of the present invention to provide a system and method of compaction and reconstruction of wavelet transform data which preserves areas of high information content in an input signal. The input signal can be a one-dimensional signal, a two-dimensional image or a many dimensional signal.

There is therefore provided, in accordance with a preferred embodiment of the present invention, a system for compaction and reconstruction of wavelet transform data of an input signal where the wavelet transform data is formed of a plurality of resolution levels. The system includes a compaction unit and a converter. The compaction unit is operative on forests of trees, wherein each tree has branches and each branch relates one datapoint in a low resolution level with the entirety of its descendant datapoints in higher resolution levels. The compaction unit prunes branches from the trees and trees from the forest to remove all but areas of high information content in the input signal, thereby to produce pruned trees. The converter converts the pruned trees to reconstructed wavelet transform data.

Additionally, in accordance with a preferred embodiment of the present invention, the compaction unit includes a) a non-null pruning unit for removing branches of the trees having datapoints thereon whose wavelet transform values are below predetermined threshold levels, producing thereby non-null pruned trees and b) a zero-crossed pruning unit for determining locations of transients in the input signal and for creating zero-crossed pruned trees having therein only datapoints of the non-null pruned trees which are at or near the locations of transients.

Moreover, in accordance with a preferred embodiment of the present invention, the datapoints of the zero-crossed pruned trees form nodes and the compaction unit further includes apparatus for associating datapoints of the non-null pruned trees with nodes of said zero-crossed pruned trees, wherein each datapoint can be within a local area of many nodes but is associated with only one of them.

Further, in accordance with a preferred embodiment of the present invention, the system includes a nullifier, operative prior to the compaction unit, for nulling those datapoints of the wavelet transform data which are below the predetermined threshold levels.

Still further, in accordance with a preferred embodiment of the present invention, the zero-crossed pruning unit includes apparatus for determining locations of zero-

crossings in the wavelet transform data.

Finally, the method performs the operations of the elements of the system.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the drawings in which:

5 Fig. 1A is a block diagram illustration of a prior multi-resolution wavelet transform operation;

Fig. 1B is a block diagram illustration of one section of a prior art wavelet transform operation for separable filtering;

10 Fig. 1C is a block diagram illustration of two sections of a prior art wavelet transform operation for non-separable filtering;

Fig. 2 is a graphical illustration of an input signal and versions of the signal after wavelet transformation having high frequency information therein;

15 Fig. 3 is a block diagram illustration of a wavelet transform compaction and reconstruction system, constructed and operative in accordance with a preferred embodiment of the present invention;

Fig. 4 is a graphical illustration of the parent-child relationships between the signals of Fig. 2, useful in understanding the operation of the present invention;

Figs. 5A and 5B are schematic illustrations of full and pruned trees, useful in understanding the operation of the system of Fig. 3;

20 Fig. 6 is a block diagram illustration of a method of nullifying some wavelet transform data, useful in the system of Fig. 3;

Fig. 7 is a flow chart illustration of the operation of a compaction unit forming part of the system of Fig. 3;

25 Fig. 8 is a flow chart illustration of a method of producing a zero-crossed tree, forming part of the operations of the compaction unit of Fig. 7;

Fig. 9 is a flow chart illustration of a method of computing zero-crossing locations forming part of the method of Fig. 8;

Fig. 10A is a schematic illustration of a local area within a two-dimensional image;

30 Fig. 10B is a schematic illustration of a local area for an odd level wherein quincunx subsampling was performed; and

Fig. 11 is a flow chart illustration of a method of producing a fully pruned tree

forming part of the method of Fig. 7.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

The present invention will be described hereinbelow in the context of one- and
5 two-dimensional input signals, it being understood that input signals of other orders of magnitude are also included within the scope of the present invention.

Reference is now made to Fig. 3 which illustrates a wavelet transform compaction and reconstruction system, constructed and operative in accordance with a preferred embodiment of the present invention. The system of the present invention typically comprises a compaction system 20 for compacting wavelet transform signals of a digital input signal and a reconstruction system 22 for reconstructing the input signal from the compacted 10 wavelet transform signals.

The compaction system 20 typically comprises a wavelet transformer 24, a nullifier 26 for zeroing out values of the wavelet transform data W which are below a predetermined human-defined threshold and a compaction unit 28 which compacts the output 15 W^0 of the nullifier 26.

The wavelet transformer 24 comprises a predetermined plurality of wavelet transform units 10 combined as described in the Background of the Present Invention. For most of the examples herein, wavelet transformer 24 comprises three wavelet transform units 20 10, however it will be understood that other numbers of units 10, typically in the range of three to eight, are suitable.

The wavelet transform units 10 can be separable or non-separable. It is noted that the type of wavelet transform utilized (either separable or non-separable) affects the percentage of compaction to be achieved, as will be described hereinbelow.

As is known in the art, the wavelet transformer 24 produces two types of signals, the core signal and the detailed signal. The core signal is the low frequency signal, labeled L3 in Fig. 1, which has a low resolution. The detailed signal comprises the high frequency signals H1 - H3 each having a different resolution.

If separable filtering was utilized on a two-dimensional input image, each high 30 frequency signal H1 - H3 is formed of three signals. For the remainder of this discussion, the term H1 will be utilized and it will refer to the output of both separable and non-separable filtering unless otherwise specified.

The compaction unit 28 typically determine the areas of "high information content" or of detail in the nullified wavelet transform data W^0 . These areas can be defined as those areas which have transients in them and which exist in most or all of the high frequency signals of the nullified wavelet data W^0 .

5 In accordance with the present invention and as illustrated in Figs. 4, 5A and 5B to which reference is now briefly made, the compaction unit 28 organizes the detailed signal as a "forest" of "trees" connecting the datapoints of the higher resolution signals (H1 and H2) with their corresponding datapoints in the lowest resolution signal (H3). The datapoints of the lowest resolution level form the "roots" of the trees and therefore, the lowest 10 resolution signal H3 is herein denoted the "root signal". There is no more than one tree for each datapoint in the root signal.

The compaction unit 28 then "prunes" the trees to remove datapoints having little information content. The compaction unit 28 can also remove trees from the forest if there is no information in a tree.

15 Figs. 5A and 5B illustrate two full trees 30 and two pruned trees 32 for two datapoints (i) and (i+1) of the root signal of a one-dimensional signal. In this example, there are five levels and thus, five high frequency signals, H1 - H5, are illustrated. In this example, each wavelet transform unit 10 reduced the resolution of the signal by 2. Therefore, each 20 datapoint at one resolution level corresponds to two datapoints in the next highest resolution level.

It is noted that each tree 30 and 32 comprises a multiplicity of branches 34 connecting each datapoint or node 35 of a signal with its corresponding datapoints of the signal having the next highest resolution. It is also noted that there are no branches 34 of the 25 pruned trees 32 which are not connected to the structure. Thus, analogous to pruning a physical tree, when one branch 34 is removed, all of the branches 34 connected to it in the higher resolution signals must also be removed.

Referring back to Fig. 3, the compaction unit 28 produces a forest of pruned trees 32, at most one tree for each datapoint in the root signal. The forest of pruned trees 32 contains generally all the essential information in the input signal, but in a more compact 30 format than the output of the wavelet transform. The compact forest of pruned trees 32 can then be transmitted as desired and as illustrated in Fig. 3. For the purposes of the present discussion, "transmission" indicates transmission between any two elements, whether through

communication devices or by storing on a storage device.

The reconstruction system 22 of the present invention typically comprises a converter 40 for converting the pruned forest back to a wavelet transform representation W^* and an inverse wavelet transformer 42 for performing the inverse wavelet transform operation, thereby to reconstruct the input signal from the wavelet transform representation W^* . Because the compaction system 20 maintains generally the essential details of the input signal, the reconstructed signal strongly resembles the input signal.

It will be appreciated that the system of the present invention can be utilized in any situation where the input signal must be reduced in size while the information content within it be maintained. For example, if the input signal is a high quality digital image, comprised of more than one-quarter million picture elements or pixels, it can be compacted with the present system for storage and reconstructed for operation. Alternatively, it can be transmitted to a second location after compaction and it can be reconstructed at the second location.

The compacted form of the wavelet transform lends itself easily to further compression, as follows: The core signal may be further compressed by known Discrete Cosine Transform (DCT) methods; the full-pruned forest topology may be encoded by standard tree encoding methods; and the datapoints of each node of the fully-pruned forest may be considered as a vector and thus, vector quantization may be applied to compress them.

Reference is now made to Fig. 6 which illustrates the operations of nullifier 26. The equations performed by the nullifier 26, for a two dimensional input signal (e.g. an image) are listed hereinbelow in the section labeled "Nullifier Equations". The index ℓ indicates the resolution level and (i,j) are spatial indices. The nullifier 26 determines a threshold level for nullifying, based on the type of input signal and on its particular features. Any values of the wavelet data less than the threshold value are nullified (e.g. set to 0).

Because the wavelet transform data W is multi-resolutonal and the values in a signal at one resolution level are not equivalent to those in a signal at another resolution level due to human response mechanisms to the specific type of input signal (audio, video, etc), the data W , in step 52, has to be normalized according to human responses.

The human response function $R(\ell,i,j)$ depends on the resolution level ℓ and also on the position of each datapoint (i,j) within the image. For the type of wavelet transforms

described hereinabove, an exemplary human response function $R(\ell, i, j)$ is:

$$R(\ell, i, j) = R(\ell) = 0.95 * 2.2^{-\ell^2} \quad (1)$$

5 Prior to normalization, in step 50, the absolute values of the data are determined, producing signal W_1 . The output of step 52 is a signal W_2 .

After normalization, the data is quantized (step 54), with a quantization level Q , to reduce the number of possible values, and then clipped (step 56) to maintain the number of possible values to a finite and reasonable amount, such as 1024.

10 It will be appreciated that steps 50 - 56 can be performed at once, if desired.

A histogram of the resultant normalized, quantized data from all of the high frequency signals, is produced in step 58. In step 60 an overall threshold level T_o is determined, based on the desired percentage, such as 80%, of the histogram to be nullified. The overall threshold level is the value of the histogram for which the desired percentage of the data has values less than or equal to it.

15 The overall threshold level cannot be utilized for the threshold operation (step 64) since its value is a normalized one. Therefore, in step 62, the threshold is "denormalized" (i.e. the inverse of the normalization operation) to provide threshold values $T(\ell, i, j)$ for each resolution level ℓ and across the image.

20 In step 64, the threshold operation is performed wherein each datapoint in each high frequency signal $H1 - H3$ of the wavelet transform data W is compared to its corresponding threshold and those datapoints which are less in absolute value than the corresponding threshold are nullified.

25 Reference is now made to Fig. 7 which illustrates the operations of compaction unit 28. Compaction unit 28 performs two types of pruning, a first pruning (step 70) to remove those branches which have null data on them, and a second pruning (step 72) to remove those branches which do not have transient data. Since transients are typically more than one datapoint wide, compaction unit 28 then enlarges (step 74) the nodes of the resultant pruned tree to include at least some of the non-null datapoints near the datapoint of the node.

30 The pruning steps 70 - 74 can be performed by creating lists of the nodes which belong in each pruned tree. In accordance with an alternative embodiment of the present invention, each datapoint can have many attributes and pruning is performed by

updating the attributes of each node which belongs in the pruned tree. The steps 70 - 74 will be described as though new pruned trees are produced, it being understood that either method can be performed.

5 Prior to discussing the specific details of each operation of compaction unit 28, we define a parent node and its child nodes of the level below it. For a one-dimensional signal, a parent node $W^0(\ell,i)$ has the following children nodes: $W^0(\ell-1,2i+1)$, $W^0(\ell-1,2i)$.

10 For a two-dimensional signal upon which separable filtering was performed, a parent node $W^0(\ell,d,i,j)$ has the following child nodes: $W^0(\ell-1,d,2i+1,2j+1)$, $W^0(\ell-1,d,2i+1,2j)$, $W^0(\ell-1,d,2i,2j+1)$ and $W^0(\ell-1,d,2i,2j)$, where d indicates the type of filtering performed. For example, $d = 1$ indicates that the rows were low pass filtered and the columns were high pass filtered, $d = 2$ indicates that both rows and columns were high pass filtered and $d = 3$ indicates that the rows were high pass filtered and the columns were low pass filtered.

15 For a two-dimensional signal upon which quincunx subsampling was performed, a parent node $W^0(\ell,i,j)$ of an even-numbered level has the following child nodes of the odd-numbered level below it: $W^0(\ell-1,2i,2j)$, $W^0(\ell-1,2i+1,2j+1)$. A parent node $W^0(\ell,i,j)$ of an odd level has the following child nodes of the even-numbered level below it: $W^0(\ell-1,i,j)$ and $W^0(\ell-1,i+1,j)$ for j even and $W^0(\ell-1,i,j)$ and $W^0(\ell-1,i-1,j)$ for j odd. It is noted that in the odd levels, if $i+j$ is odd, the datapoint does not exist.

20 Similar relationships can be built for input signals of higher dimensions.

In step 70, the compaction unit 28 reviews the entirety of nodes in the forest of trees. For each node which is null and whose descendants do not have significant information in them, where "significance" is defined hereinbelow in equation (2), the node and its descendants are removed from the tree to which it belongs. Typically, the compaction unit 28 begins the review with the root signal H3.

25 The significance S of the data in the descendants of the node is determined as follows:

$$S = \sum W^0(*)R(*) \quad (2)$$

30 where (*) indicates indices and denotes a descendant of the node 35, $W^0(*)$ is the value of the descendant, $R(*)$ is the value of the normalization function for the descendant, and the

sum is over all the descendants of the node. S has to be above a predetermined threshold value in order for the node to be kept within the pruned tree structure.

For those nodes 35 which have null values, the compaction unit 28 determines whether or not the branches 34 attached to the node have significant data in them. If so, it saves the node as part of the pruned tree structure.

The compaction unit 28 repeats the above determination for each of the datapoints in the root signal and for each of the high frequency signals. The result is a forest of "non-null" pruned trees.

In step 72, the compaction unit 28 reviews the non-null pruned tree and the wavelet transformation W to determine the areas corresponding to "transients". A transient is herein defined as the location where the wavelet transformation value changes from positive to negative or vice-versa; in other words, a transient is the location of a "zero-crossing". By choosing the appropriate wavelet bases for the wavelet transformation operation (performed by wavelet transformer 24 (Fig. 3)), the zero crossings can be forced to be at locations of maximum transitions in the input signal.

Selection of appropriate wavelet bases is discussed in the following article which is incorporated herein by reference: Stephane Mallet, "Zero-Crossings of a Wavelet Transform", IEEE Transactions on Information Theory Vol. 37, No. 4, July 1991, pp. 1020 - 1033.

The operations performed are provided in the flow charts of Figs. 8 and 9, to which reference is now made. Fig. 8 illustrates the operations performed to produce the zero-crossed pruned trees and Fig. 9 illustrates the operations performed to determine where the zero-crossings occur.

In step 80 of Fig. 8, which is detailed in Fig. 9, the compaction unit 28 determines which nodes 35 of the non-null pruned forest are part of transients. The location of the transients are determined through an iterative process operating on the non-nullified wavelet transform W.

The iteration process shown in Fig. 9 operates on a "local area" or neighborhood in W of each node 35 in the non-null pruned tree. Local areas for a two-dimensional input signal are shown in Figs. 10A and 10B to which reference is now briefly made. For signals upon which two-dimensional separable sampling was performed, the neighborhood is typically an NxN square of data surrounding the current node 35, where N

is typically 3 - 7. This is illustrated in Fig. 10A. If N is odd, then the node 35 sits in the very center of the neighborhood. Otherwise, it sits slightly to one side. If two-dimensional quincunx subsampling was performed, for even numbered levels, the local area is as shown in Fig. 10A. For odd numbered levels, the neighborhood is an NxN diamond, as shown in Fig. 10B.

Referring back to Figs. 8 and 9, in step 82, the compaction unit 28 selects a first node 35 of the non-null pruned forest as the current "nucleus". In step 86 a zero-crossing location for the current local area surrounding the current nucleus is computed. If the location of the zero-crossing happens to be at the current nucleus (step 88), then the current nucleus is "zero-crossed" and there are no more iterations. Otherwise, the current nucleus is set to be at the location of the zero-crossing and the process is repeated for the new local area around the current nucleus. The iterations continue until the current nucleus is at the location of a zero-crossing or a maximum number of iterations (typically 5 to 10) has been reached. The last value for the zero-crossing, no matter how far from the original node, is saved (step 94) and the process repeated for each node of the non-null pruned forest.

The equations for determining the location of the zero-crossing are provided hereinbelow in the section labeled "Zero Crossing Equations".

The remaining steps of Fig. 8 create "zero-crossed" trees which contain only those nodes which are zero-crossed or which have a zero-crossing within their local area. Specifically, for each non-null pruned tree, the compaction unit 28 scans the first resolution signal H1. Whenever (step 102) it finds a "zero-crossed node" (i.e. one on which the zero-crossing occurs directly) from the list produced in step 90, it checks if the node is already a node in some zero-crossed tree (step 104) in which case, it stops the present determinations and adds a temporary list STL to the zero-crossed tree to which the node belongs (step 110). If it is not already part of the zero-crossed tree, the unit 28 saves the node in the temporary list STL (step 106).

Unit 28 then sets the current node to its parent node (step 108) and then has to determine if the parent node should be added to the zero-crossed tree. It repeats steps 104 and 106 and, if the parent is not a root, the following criteria are checked: (step 112) is the parent node zero-crossed, (step 114) does it have a zero-crossing in its local area based on the information from the method of Fig. 9; or (step 116) are its parent and its child both zero-crossed. In other words, step 116 indicates that the parent is a small gap in a branch of zero-

crossed nodes and is a relaxation criterion vis-a-vis the other criteria.

If the parent node fulfills one of the criteria, it is added to the temporary list STL and the next parent is taken. The determination is repeated until a datapoint from the root signal is reached (step 118) or when a node which is already part of a zero-crossed tree is found. Otherwise, the determination for this node of the first high frequency signal level H1 is finished and the next node is taken. When the nodes of H1 are exhausted, the process is repeated for the other high frequency signals, starting from level H2.

The nodes are scanned in quanta of local area size. Thus, the first resolution level H1 is divided into units the size of one local area. For a two-dimensional input image, the units are squares or diamonds for the odd levels from quincunx subsampling. During a first pass, the datapoints in the same position in each local area, for example in the top left corner, are checked. In the next pass, the datapoints from the next position, for example, the next position in the top row, are checked.

The datapoints are first checked that they are nodes of the non-null tree. If not, the next datapoint (i.e. from the next local area) is taken.

The result of the operations of Fig. 8 are zero-crossed trees having therein only those datapoints which are at or near transients. However, a human typically views a transient as having a certain thickness. Thus, the zero-crossed trees do not contain enough information. They must be filled with those nodes of the non-null pruned trees which are near the nodes of the zero-crossed trees.

Therefore, in step 74 of Fig. 7, the compaction unit 28 adds data to the nodes of the zero-crossed trees. The operations performed in step 74 are listed, in flow-chart format, in Fig. 11.

The nodes of the new trees, called herein "fully pruned trees", are first set to be the nodes of the zero-crossed trees (step 130).

The compaction unit 28 then scans the nodes N of each non-null pruned tree to determine which nodes are either already part of the structure (step 132) or are part of the local area of at least one node of the corresponding zero-crossed tree (step 134). A list M is produced of the nodes of the zero-crossed trees whose local area the current non-null node N is in.

If a node is neither part of the structure or part of the local area of a node, a new node is tested (step 136). If M has one member, m*, the current node N is added to the

fully pruned tree at node m^* (step 138). If M has many members (i.e. the node N is part of many local areas), the node m which is closest to the current node N is selected as m^* (step 140) and current node N is added to the fully pruned tree at node m^* (step 138). The process is then repeated for all nodes in each non-null pruned tree.

5 The process of Fig. 11 ensures that, if a node of a non-null tree should be included in the fully pruned tree, it is included at only one node. Thus, the fully pruned tree contains the areas of high information content in a non-redundant manner.

10 The converter 40 (Fig. 3) of the decompression unit 22 reproduces the wavelet transform data W^* by placing the value of zero at every pixel not included in the list. The inverse wavelet transformer 42 then reproduces the input signal from the uncompacted wavelet transform data W^* .

Nullifier Equations:

$$W_1(l,i,j) = |W(l,i,j)| \quad (3)$$

$$15 \quad W_2(l,i,j) = W_1(l,i,j) + R(l,i,j) \quad (4)$$

$$W_3(l,i,j) = \lfloor \frac{W_2(l,i,j)}{Q} + 0.5 \rfloor \quad (5)$$

where $\lfloor x \rfloor$ is the floor of X which is the maximum integer smaller than X .

$$W_4(l,i,j) = MIN[W_3(l,i,j), W_{MAX}] \quad (6)$$

$$20 \quad T_o = [K | \sum_{k=0}^K HST(k) \leq \frac{1}{P} \sum_{k=0}^{W_{MAX}} HST(k), \sum_{k=0}^{K+1} HST(k) > \frac{1}{P} \sum_{k=0}^{W_{MAX}} HST(k) |] \quad (7)$$

$$T(l,i,j) = \frac{QT_o}{R(l,i,j)} \quad (8)$$

$$W^o(l,i,j) = \begin{cases} 0 & \text{if } W(l,i,j) \leq T(l,i,j) \\ W(l,i,j) & \text{otherwise} \end{cases} \quad (9)$$

Where $R(l,i,j)$ is a function of "importance" which may be fixed for each level

($R(l,i,j)=R(l)$) or data dependent (for example, the human visual response function may

5 serve as the function R for image data), Q is a predetermined quantization step and W_{MAX}

is a predetermined maximum number of bins in the histogram.

Zero Crossing Equations:

The following is noted about the zero crossing equations: a) they are written for a two-dimensional input signal, b) as the zero crossing is done within the level, we drop the level index, c) $W(i,j)$ denotes a detailed signal datapoint and d) the operations are done for a w
10 within a local area nucleus.

$$S_{POS} = \sum_{w>0} W(i,j) \quad (10)$$

$$15 \quad \overline{X_{POS}} = \sum_{w>0} i \circ W(i,j) \quad (11)$$

$$\overline{Y_{POS}} = \sum_{w>0} j \circ W(i,j) \quad (12)$$

$$S_{neg} = - \sum_{w<0} W(i,j) \quad (13)$$

$$\overline{X_{neg}} = - \sum_{w<0} i \circ W(i,j) \quad (14)$$

$$\overline{Y_{neg}} = - \sum_{w>0} j \circ W(i_p, j) \quad (15)$$

$$X_{zc} = \frac{\overline{X_{neg}} S_{pos} + \overline{X_{pos}} S_{neg}}{S_{pos} + S_{neg}} \quad (16)$$

$$Y_{zc} = \frac{\overline{Y_{neg}} S_{pos} + \overline{Y_{pos}} S_{neg}}{S_{pos} + S_{neg}} \quad (17)$$

5

It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather the scope of the present invention is defined by the claims which follow:

CLAIMS

1. A system for compaction and reconstruction of wavelet transform data of an input signal, the wavelet transform data formed of a plurality of resolution levels, the system comprising:

a compaction unit, operative on forests of trees, each tree having branches which relate one datapoint in a low resolution level with the entirety of its descendant datapoints in higher resolution levels, for pruning branches from said trees and trees from said forest to remove all but areas of high information content in said input signal and for producing thereby pruned trees; and

10 a converter for converting said pruned trees to reconstructed wavelet transform data.

2. A system according to claim 1 wherein said compaction unit comprises a non-null pruning unit for removing branches of said trees having datapoints thereon whose wavelet transform values are below predetermined threshold levels and for producing thereby non-null pruned trees.

3. A system according to claim 2 and wherein said compaction unit additionally comprises a zero-crossing pruning unit for determining locations of transients in said input signal and for creating zero-crossed pruned trees having therein only datapoints of said non-null pruned trees which are at or near said locations of transients.

4. A system according to claim 3 wherein said datapoints of said zero-crossing pruned trees form nodes and wherein said compaction unit further comprises means for associating datapoints of said non-null pruned trees with nodes of said zero-crossed pruned trees, wherein each datapoint can be within a local area of many nodes but is associated with only one of them.

25 5. A system according to claim 2 and including a nullifier, operative prior to said compaction unit, for nulling those datapoints of said wavelet transform data which are below said predetermined threshold levels.

6. A system according to claim 3 and wherein said zero-crossed pruning unit includes means for determining locations of zero-crossings in said wavelet transform data.

5 7. A method of compaction and reconstruction of wavelet transform data of an input signal, the wavelet transform data formed of a plurality of resolution levels, the method comprising the steps of:

10 pruning branches from trees and trees from forests to remove all but areas of high information content in said input signal and producing thereby pruned trees, wherein each tree has branches and each branch relates one datapoint in a low resolution level with the entirety of its descendant datapoints in higher resolution levels; and

converting said pruned trees to reconstructed wavelet transform data.

15 8. A method according to claim 7 wherein said step of pruning comprises the step of removing those branches of said trees having datapoints thereon whose wavelet transform values are below predetermined threshold levels and producing thereby non-null pruned trees

20 9. A method according to claim 8 and also comprising the step of determining locations of transients in said input signal and creating zero-crossed pruned trees having therein only datapoints of said non-null pruned trees which are at or near said locations of transients.

25 10. A method according to claim 8 wherein said datapoints of said zero-crossed pruned trees form nodes and wherein said step of pruning further comprises the step of associating datapoints of said non-null pruned trees with nodes of said zero-crossed pruned trees, wherein each datapoint can be within a local area of many nodes but is associated with only one of them.

30 11. A method according to claim 8 and including the step, operative prior to said compaction unit, of nulling those datapoints of said wavelet transform data which are below said predetermined threshold levels.

12. A method according to claim 8 and wherein said step of determining includes the step of determining locations of zero-crossings in said wavelet transform data.

• Relevant Technical Fields	Search Examiner M K REES
(i) UK Cl (Ed.N) H4F (FRT, FRW, FRX)	Date of completion of Search 13 JANUARY 1995
(ii) Int Cl (Ed.6) H04N (7/26)	Documents considered relevant following a search in respect of Claims :- 1 TO 12
(ii) ONLINE DATABASE: WPI	

Categories of documents

- X: Document indicating lack of novelty or of inventive step.
- Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.
- A: Document indicating technological background and/or state of the art.
- P: Document published on or after the declared priority date but before the filing date of the present application.
- E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.
- &: Member of the same patent family; corresponding document.

Category	Identity of document and relevant passages		Relevant to claim(s)
X, P	WO 94/23385 A2	(AS LEWIS & GP KNOWLES) 13 October 1994; see Figures 23 to 29; page 48, line 14 to page 59, line 13	1, 2, 7, 8
X	WO 93/17524 A1	(GENERAL ELECTRIC CO) see Figures 3, 9 and 10; page 10, line 9 to page 21, line 2	1, 2, 7, 8

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).